Patents Office
Government Buildings
Hebron Road
Kilkenny

JC978 U.S. PTO
10/091647
03/05/02

#2
Priority Paper
MAR
5/8/02

I HEREBY CERTIFY that annexed hereto is a true copy of documents filed in connection with the following patent application:

Application No.	S2001/0223
Date of Filing	8 March 2001
Applicant	RICHMOUNT COMPUTERS LIMITED, an Irish Company of Maple House, South County Business Park, Leopardstown, Dublin 18, Ireland.

Dated this 2 day of October 2001.

An officer authorised by the
Controller of Patents, Designs and Trademarks.

S010223

FORM NO. 1

REQUEST FOR THE GRANT OF A PATENT
PATENTS ACT, 1992

The Applicant named herein hereby request

__ the grant of a patent under Part II of the Act

_X_ the grant of a short-term patent under Part III of the Act

on the basis of the information furnished hereunder.

1. **APPLICANT**

Name	**Richmount Computers Limited**
Address	**Maple House, South County Business Park, Leopardstown, Dublin 18, Ireland.**
Description/Nationality	**An Irish Company**

2. **TITLE OF INVENTION**

"Reset Facility for Redundant Processor Using only Fibre Channel Loop"

3. **DECLARATION OF PRIORITY ON BASIS OF PREVIOUSLY FILED APPLICATION FOR SAME INVENTION (SECTIONS 25 & 26)**

Previous filing date	Country in or for which filed	Filing No.

4. **IDENTIFICATION OF INVENTOR(S)**
Name(s) of person(s) believed by Applicant(s) to be the inventor(s)
1. Aedan Diarmuid Coffey

Address
1. 1 Sion Hermitage, Sion Road, County Kilkenny, Ireland.

5. **STATEMENT OF RIGHT TO BE GRANTED A PATENT (SECTION 17(2)(B))**

By virtue of

Contd./...

6. **ITEMS ACCOMPANYING THIS REQUEST - TICK AS APPROPRIATE**

(i) X prescribed filing fee EUR 63.49 (IR£50.00)

(ii) _ specification containing a description and claims

X specification containing a description only

X Drawings referred to in description or claims

(iii) _ An abstract

iv) _ Copy of previous application(s) whose priority is claimed

(v) _ Translation of previous application whose priority is claimed

(vi) _ Authorisation of Agent (this may be given at 8 below if this Request is signed by the Applicant(s))

7. **DIVISIONAL APPLICATION**

The following information is applicable to the present application which is made under Section 24

Earlier Application No:
Filing Date:

8. **AGENT**

The following is authorised to act as agent in all proceedings connected with the obtaining of a Patent to which this request relates and in relation to any patent granted -

Name	**Address**
F. R. KELLY & CO.	at their address as recorded for the time being in the Register of Patent Agents

9. **ADDRESS FOR SERVICE (IF DIFFERENT FROM THAT AT 8)**

RICHMOUNT COMPUTERS LIMITED
F. R. KELLY & CO.

By: _____
EXECUTIVE

Date: March 8, 2001

Oifig na bPaitinní
FAIGHTE
8 MAR 2001

TRUE COPY AS LODGED

APPLICATION No. S010223

Reset facility for redundant processor using only fibre channel loop

The present invention relates to a method for controlling and resetting one processor in a redundant pair using only the fibre channel loops.

Any highly available system with dual (or more) redundant processors requires a method of forcing a faulty processor off the system, to prevent it disrupting normal operation. Normally this functionality requires dedicated cabling, thus making it difficult to expand the system easily beyond a dedicated backplane.

Redundant processors would normally have dedicated wiring between them to allow an errant one to be reset or powered down.

Dedicated wiring requires extra PWB traces and extra cabling between processors. This is both expensive and contributes to unreliability.

This invention overcomes the above problem by using special purpose hardware, referred to as a HASC (High Availability Support Chip). Using a HASC, one processor can interrogate and control the reset signals of another, thus forcing it off the fibre channel loop if necessary.

This invention could be used in a high availability version of Intelligent Network Application Protocol

INAP, or any other redundant processing system using fibre channel as a communications medium.

This invention could allow high availability server systems to be offered using existing backplanes and cabling systems.

The invention allows the building of a high availability, scaleable file server that does not require additional inter-processor wiring.

Furthermore, the invention improves performance by accelerating lock management functions.

Other features include:

High availability: No single point of failure.
Scalable: More INAP's can be added allowing an approximately linear performance increase.
Fast, redundant, distributed, scalable lock management.
Ability for one INAP to reset another if it detects that it is faulty.
All communications over FC-AL loop, thus scalable beyond a shelf, even into two separate geographical locations.

Details:

Each INAP has a FC-AL sub-system as shown in Figure 1. Note that the FC-AL interface chips are already on the existing INAP design.

CAM: Content Addressable Memory

A special type of memory with a built-in search capability. You give it the contents you are looking for, it returns the address if there is a match. Commonly used in routers and switches.

SERDES: Serialiser/Deserialiser

Converts serial FC data into parallel data at 1/10th or 1/20th of the speed.

HASC: High Availability Support Chip

Overview Points:

1. HASC allows the CAM to be read and written by the INAP CPU. The CAM is used to store lock data.
2. The CAM can be searched by the local CPU or by any other device on the FC-AL loop, via the SERDES and HASC.
3. The HASC has the facility to reset or interrupt the INAP module, via a command received over either FC-AC loop.
4. The HASC is 100% hardware, thus allowing fast CAM lookups and predictable reset operation.

OVERALL SYSTEM OPERATION

A system comprises one or more shelves, with two or more INAP-HAs.

Each INAP-HA has two expansion ports, thus allowing expansion of shelves containing IAP-HAs only, no I/O cards.

At system initialisation time each INAP-HA twins with a "buddy", preferably not in the same shelf (for added reliability).

Each INAP-HA regularly talks with its buddy, checking it's OK, using a high level watchdog type system, over the FC-AC loops.

Whenever an INAP-HA takes out a lock on a file system resource its buddy also put the lock data into its CAM. Thus the lock data is redundantly stored without every INAP-HA having to hold all the lock data. The solution is thus scalable.

When an INAP-HA wishes to check a lock it puts a broadcast frame on the loop (using either the standard if chip or the HASC, TBD). Each HASC retrieves the frame, searches it's CAM and passes the frame on, marking it as a "hit" if the lock exists. The frame will arrive back at the originator having been checked by all HASCs.

If one INAP-HA detects that it's "buddy" has not sent "I'm OK" type messages for an extended time it first checks the loop is up, then if more attempts to communicate fail it sends a "reset" frame to the HASC on it's buddy. This feature is called STOMITH (Shoot The Other Machine In The Head). One of the INAP-HAs in a set is designated the master, it's watchdog timeout

is set to a lesser time than the other, in order to prevent both INAP-HAs resetting each other at the same time.

During normal operation the INAP-HA's load share, in an active-active manner.

If an INAP-HA loses its buddy it can buddy up with a spare, if available.

There may be a requirement for more INAP-HA processors than the natural limit (approx 8 shelves, 16 INAP-HA's). In this case there are four alternatives.

- Add extra shelves with no drives.
- Re-package INAP-HA into a format that can be loaded from the front, instead of one or more disks, using the SCA connectors.
- Design a custom backplane, capable of taking lots of INAP-HAs, in a front loadable format.
- Design metalwork capable of holding INAP-HAs.

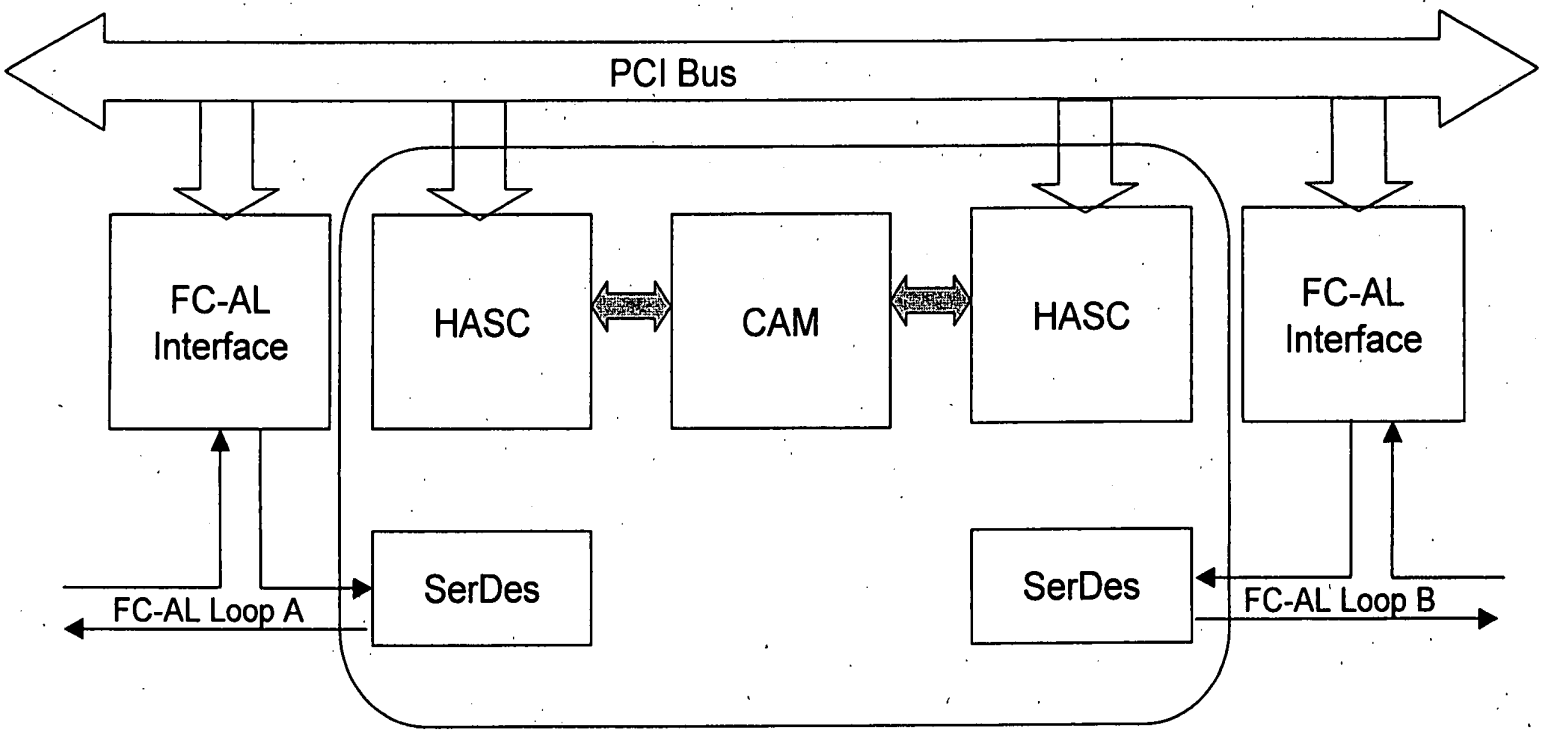
PCI Bus
FC-AL Interface
HASC
CAM
HASC
FC-AL Interface
SerDes
SerDes
FC-AL Loop A
FC-AL Loop B

Figure 1